

A tutorial for the sam command language

Rob Pike

ABSTRACT

sam is an interactive text editor with a command language that makes heavy use of regular expressions. Although the language is syntactically similar to **ed**(1), the details are interestingly different. This tutorial introduces the command language, but does not discuss the screen and mouse interface. With apologies to those unfamiliar with the Ninth Edition Blit software, it is assumed that the similarity of **sam** to **mux**(9) at this level makes **sam**'s mouse language easy to learn.

The **sam** command language applies identically to two environments: when running **sam** on an ordinary terminal (*via* **sam -d**), and in the command window of a *downloaded sam*, that is, one using the bitmap display and mouse.

Introduction

This tutorial describes the command language of **sam**, an interactive text editor that runs on Blits and some computers with bitmap displays. For most editing tasks, the mouse-based editing features are sufficient, and they are easy to use and to learn.

The command language is often useful, however, particularly when making global changes. Unlike the commands in **ed**, which are necessary to make changes, **sam** commands tend to be used only for complicated or repetitive editing tasks. It is in these more involved uses that the differences between **sam** and other text editors are most evident.

sam's language makes it easy to do some things that other editors, including programs like **sed** and **awk**, do not handle gracefully, so this tutorial serves partly as a lesson in **sam**'s manner of manipulating text. The examples below therefore concentrate entirely on the language, assuming that facility with the use of the mouse in **sam** is at worst easy to pick up. In fact, **sam** can be run without the mouse at all (not *downloaded*), by specifying the **-d** flag, and it is this domain that the tutorial occupies; the command language in these modes are identical.

A word to the Unix adept: although **sam** is syntactically very similar to **ed**, it is fundamentally and deliberately different in design and detailed semantics. You might use knowledge of **ed** to predict how the substitute command works, but you'd only be right if you had used some understanding of **sam**'s workings to influence your prediction. Be particularly careful about idioms. Idioms form in curious nooks of languages and depend on undependable peculiarities. **ed** idioms simply don't work in **sam**: **1,\$s/a/b/** makes one substitution in the whole file, not one per line. **sam** has its own idioms. Much of the purpose of this tutorial is to publish them and make fluency in **sam** a matter of learning, not cunning.

The tutorial depends on familiarity with regular expressions, although some experience with a more traditional Unix editor may be helpful. To aid readers familiar with **ed**, I have pointed out in square brackets [] some of the relevant differences between **ed** and **sam**. Read these comments only if you wish to understand the differences; the lesson is about **sam**, not **sam** vs. **ed**. Another typographic convention is that output appears in **this font**, while typed input appears as *slanty text*.

Nomenclature: **sam** keeps a copy of the text it is editing. This copy is called a *file*. To avoid confusion, I have called the permanent storage on disc a *Unix file*.

Text

To get started, we need some text to play with. Any text will do; try something from James Gosling's Emacs manual:

\$ sam -d

a

This manual is organized in a rather haphazard manner. The first several sections were written hastily in an attempt to provide a general introduction to the commands in Emacs and to try to show the method in the madness that is the Emacs command structure.

Sam -d starts **sam** running. The **a** command adds text until a line containing just a period, and sets the *current text* (also called *dot*) to what was typed — everything between the **a** and the period. [**ed** would leave dot set to only the last line.] The **p** command prints the current text:

p

This manual is organized in a rather haphazard manner. The first several sections were written hastily in an attempt to provide a general introduction to the commands in Emacs and to try to show the method in the madness that is the Emacs command structure.

[Again, **ed** would print only the last line.] The **a** command adds its text *after* dot; the **i** command is like **a**, but adds the text *before* dot.

i

Introduction

.

p

Introduction

There is also a **c** command that changes (replaces) the current text, and **d** that deletes it; these are illustrated below.

To see all the text, we can specify what text to print; for the moment, suffice it to say that *O,\$* specifies the entire file. [**ed** users would probably type *1,\$*, which in practice is the same thing, but see below.]

O,\$p

Introduction

This manual is organized in a rather haphazard manner. The first several sections were written hastily in an attempt to provide a general introduction to the commands in Emacs and to try to show the method in the madness that is the Emacs command structure.

Except for the **w** command described below, *all* commands, including **p**, set dot to the text they touch. Thus, **a** and **i** set dot to the new text, **p** to the text printed, and so on. Similarly, all commands (except **w**) by default operate on the current text [unlike **ed**, for which some commands (such as **g**) default to the entire file].

Things are not going to get very interesting until we can set dot arbitrarily. This is done by *addresses*, which specify a piece of the file. The address **1**, for example, sets dot to the first line of the file.

1p

Introduction

c

Preamble

The **c** command didn't need to specify dot; the **p** left it on line one. It's therefore easy to delete the

first line utterly; the last command left dot set to line one:

```
d
1p
```

This manual is organized in a rather haphazard manner. The first

(Line numbers change to reflect changes to the file.)

The address `/text/` sets dot to the first appearance of `text`, after dot. [`ed` matches the first line containing `text`.] If `text` is not found, the search restarts at the beginning of the file and continues until dot.

```
/Emacs/p
Emacs
```

It's difficult to indicate typographically, but in this example no newline appears after **Emacs**: the text to be printed is the string 'Emacs', exactly. (The final **p** may be left off — it is the default command. When downloaded, however, the default is instead to select the text, to highlight it, and to make it visible by moving the window on the file if necessary. Thus, `/Emacs/` indicates on the display the next occurrence of the text.)

Imagine we wanted to change the word **haphazard** to **thoughtless**. Obviously, what's needed is another **c** command, but the method used so far to insert text includes a newline. The syntax for including text without newlines is to surround the text with slashes (which is the same as the syntax for text searches, but what is going on should be clear from context). The text must appear immediately after the **c** (or **a** or **i**). Given this, it is easy to make the required change:

```
/haphazard/c/thoughtless/
1p
```

This manual is organized in a rather thoughtless manner. The first

[Changes can always be done with a **c** command, even if the text is smaller than a line]. You'll find that this way of providing text to commands is much more common than is the multiple-lines syntax. If you want to include a slash `/` in the text, just precede it with a backslash `\`, and use a backslash to protect a backslash itself.

```
/Emacs/c/Emacs\\360/
4p
```

general introduction to the commands in Emacs\360 and to try to show

We could also make this particular change by

```
/Emacs/a\\360/
```

This is as good a place as any to introduce the **u** command, which undoes the last command. A second **u** will undo the penultimate command, and so on.

```
u
4p
```

general introduction to the commands in Emacs and to try to show

```
u
3p
```

This manual is organized in a rather haphazard manner. The first

Undoing can only back up; there is no way to undo a previous **u**.

Addresses

We've seen the simplest forms of addresses, but there is more to learn before we can get too much further. An address selects a region in the file — a substring — and therefore must define the beginning and the end of a region. Thus, the address **13** selects from the beginning of line thirteen to the end of line thirteen, and `/Emacs/` selects from the beginning of the word 'Emacs' to the end.

Addresses may be combined with a comma:

13,15

selects lines thirteen through fifteen. The definition of the comma operator is to select from the beginning of the left hand address (the beginning of line 13) to the end of the right hand address (the end of line 15).

A few special simple addresses come in handy: . (a period) represents dot, the current text, 0 (line zero) selects the null string at the beginning of the file, and \$ selects the null string at the end of the file [not the last line of the file]. Therefore,

0,13

selects from the beginning of the file to the end of line thirteen,

.,\$

selects from the beginning of the current text to the end of the file, and

0,\$

selects the whole file [that is, a single string containing the whole file, not a list of all the lines in the file].

These are all *absolute* addresses: they refer to specific places in the file. **sam** also has relative addresses, which depend on the value of dot, and in fact we have already seen one form: **/Emacs/** finds the first occurrence of **Emacs** searching forwards from dot. Which occurrence of **Emacs** it finds depends on the value of dot. What if you wanted the first occurrence *before* dot? Just precede the pattern with a minus sign, which reverses the direction of the search:

-/Emacs/

In fact, the complete syntax for forward searching is

+/Emacs/

but the plus sign is the default, and in practice is rarely used. Here is an example that includes it for clarity:

0+/Emacs/

selects the first occurrence of **Emacs** in the file; read it as “go to line 0, then search forwards for **Emacs**.” Since the **+** is optional, this can be written **0/Emacs/**. Similarly,

\$-/Emacs/

finds the last occurrence in the file, so

0/Emacs/,\$-/Emacs/

selects the text from the first to last **Emacs**, inclusive. Slightly more interesting:

/Emacs/+/Emacs/

(there is an implicit **+.+** at the beginning) selects the second **Emacs** following dot.

Line numbers may also be relative.

-2

selects the second previous line, and

+5

selects the fifth following line (here the plus sign is obligatory).

Since addresses may select (and dot may be) more than one line, we need a definition of ‘previous’ and ‘following.’ ‘previous’ means *before the beginning* of dot, and ‘following’ means *after the end*

of dot. For example, if the file contains AAAA, with dot set to the middle two A's (the slanting characters), `-/A/` sets dot to the first A, and `+/A/` sets dot to the last A. Except under odd circumstances (such as when the only occurrence of the text in the file is already the current text), the text selected by a search will be disjoint from dot.

To select the **troff -ms** paragraph containing dot, however long it is, use

```
-.PP/,.PP/-1
```

which will include the **.PP** that begins the paragraph, and exclude the one that ends it.

When typing relative line number addresses, the default number is 1, so the above could be written slightly more simply:

```
-.PP/,.PP/-
```

What does the address `+1-1` or the equivalent `+-` mean? It looks like it does nothing, but recall that dot need not be a complete line of text. `+1` selects the line after the end of the current text, and `-1` selects the line before the beginning. Therefore `+1-1` selects the line before the line after the end of dot, that is, the complete line containing the end of dot. We can use this construction to expand a selection to include a complete line, say the first line in the file containing **Emacs**:

```
0/Emacs/+-p
```

general introduction to the commands in Emacs and to try to show

The address `+-` is an idiom.

Loops

Above, we changed one occurrence of **Emacs** to **Emacs\360**, but if the name of the editor is really changing, it would be useful to change *all* instances of the name in a single command. **sam** provides a command, **x** (extract), for just that job. The syntax is `x/pattern/command`. For each occurrence of the pattern in the selected text, **x** sets dot to the occurrence and runs command. For example, to change **Emacs** to **vi**,

```
0,$x/Emacs/c/vi/
```

```
0,$p
```

This manual is organized in a rather haphazard manner. The first several sections were written hastily in an attempt to provide a general introduction to the commands in **vi** and to try to show the method in the madness that is the **vi** command structure.

This works by subdividing the current text (`0,$` — the whole file) into appearances of its textual argument (**Emacs**), and then running the command that follows (`c/vi/`) with dot set to the text. We can read this example as, “find all occurrences of **Emacs** in the file, and for each one, set the current text to the occurrence and run the command `c/vi/`, which will replace the current text by **vi**.” [This command is somewhat similar to **ed**'s **g** command. The differences will develop below, but note that the default address, as always, is dot rather than the whole file.]

A single **u** command is sufficient to undo an **x** command, regardless of how many individual changes the **x** makes.

```
u
```

```
0,$p
```

This manual is organized in a rather haphazard manner. The first several sections were written hastily in an attempt to provide a general introduction to the commands in Emacs and to try to show the method in the madness that is the Emacs command structure.

Of course, **c** is not the only command **x** can run. An **a** command can be used to put proprietary markings on **Emacs**:

```
0,$x/Emacs/a/{TM}/  
/Emacs/+-p
```

general introduction to the commands in Emacs{TM} and to try to show

[There is no way to see the changes as they happen, as in `ed`'s `g/Emacs/s//&{TM}/p`; see the section on Multiple Changes, below.]

The `p` command is also useful when driven by an `x`, but be careful that you say what you mean;

```
0,$x/Emacs/p  
EmacsEmacs
```

since `x` sets dot to the text in the slashes, printing only that text is not going to be very informative. But the command that `x` runs can contain addresses. For example, if we want to print all lines containing **Emacs**, just use `+-:`

```
0,$x/Emacs/+-p  
general introduction to the commands in Emacs{TM} and to try to show  
the method in the madness that is the Emacs{TM} command structure.
```

Finally, let's restore the state of the file with another `x` command, and make use of a handy shorthand: a comma in an address has its left side default to `0`, and its right side default to `$`, so the easy-to-type address `,` refers to the whole file:

```
,x/Emacs/ /{TM}/d  
,p  
This manual is organized in a rather haphazard manner. The first  
several sections were written hastily in an attempt to provide a  
general introduction to the commands in Emacs and to try to show  
the method in the madness that is the Emacs command structure.
```

Notice what this `x` does: for each occurrence of Emacs, find the {TM} that follows, and delete it.

The 'text' `sam` accepts for searches in addresses and in `x` commands is not simple text, but rather *regular expressions*. Unix has several distinct interpretations of regular expressions. The form used by `sam` is that of `regexp(6)`, including parentheses `()` for grouping and an 'or' operator `|` for matching strings in parallel. `sam` also matches the character sequence `\n` with a newline character. Replacement text, such as used in the `a` and `c` commands, is still plain text, but the sequence `\n` represents newline in that context, too.

Here is an example. Say we wanted to double space the document, that is, turn every newline into two newlines. The following all do the job:

```
,x/\n/ a/\n/  
,x/\n/ c/\n\n/  
,x/$/ a/\n/  
,x/^/ i/\n/
```

The last example is slightly different, because it puts a newline *before* each line; the other examples place it after. The first two examples manipulate newlines directly [something outside `ed`'s ken]; the last two use regular expressions: `$` is the empty string at the end of a line, while `^` is the empty string at the beginning.

These solutions all have a possible drawback: if there is already a blank line (that is, two consecutive newlines), they make it much larger (four consecutive newlines). A better method is to extend every group of newlines by one:

```
,x/\n+/ a/\n/
```

The regular expression operator `+` means 'one or more;' `\n+` is identical to `\n\n*`. Thus, this example takes every sequence of newlines and adds another to the end.

A more common example is indenting a block of text by a tab stop. The following all work, although the first is arguably the cleanest (the blank text in slashes is a tab):

```
,x/^/a/ /  
,x/^/c/ /  
,x/.*\n/i/ /
```

The last example uses the pattern (idiom, really) `.*\n` to match lines: `.*` matches the longest possible string of non-newline characters. Taking initial tabs away is just as easy:

```
,x/^ /d
```

In these examples I have specified an address (the whole file), but in practice commands like these are more likely to be run without an address, using the value of dot set by selecting text with the mouse.

Conditionals

The `x` command is a looping construct: for each match of a regular expression, it extracts (sets dot to) the match and runs a command. `sam` also has a conditional, `g`: `g/pattern/command` runs the command if dot contains a match of the pattern *without changing the value of dot*. The inverse, `v`, runs the command if dot does *not* contain a match of the pattern. (The letters `g` and `v` are historical and have no mnemonic significance. You might think of `g` as ‘guard.’) [ed users should read the above definitions very carefully; the `g` command in `sam` is fundamentally different from that in `ed`.] Here is an example of the difference between `x` and `g`:

```
,x/Emacs/c/vi/
```

changes each occurrence of the word **Emacs** in the file to the word **vi**, but

```
,g/Emacs/c/vi/
```

changes the *whole file* to **vi** if there is the word **Emacs** anywhere in the file.

Neither of these commands is particularly interesting in isolation, but they are valuable when combined with `x` and with themselves.

Composition

One way to think about the `x` command is that, given a selection (a value of dot) it iterates through interesting subselections (values of dot within). In other words, it takes a piece of text and cuts it into smaller pieces. But the text that it cuts up may already be a piece cut by a previous `x` command or selected by a `g`. `sam`’s most interesting property is the ability to define a sequence of commands to perform a particular task.[†] A simple example is to change all occurrences of **Emacs** to **emacs**; certainly the command

```
,x/Emacs/ c/emacs/
```

will work, but we can use an `x` command to save retyping most of the word **Emacs**:

```
,x/Emacs/ x/E/ c/e/
```

(Blanks can be used to separate commands on a line to make them easier to read.) What this command does is find all occurrences of **Emacs** (`x/Emacs/`), and then *with dot set to that text*, find all occurrences of the letter **E** (`x/E/`), and then *with dot set to that text*, run the command `c/e/` to change the character to lower case. Note that the address for the command — the whole file, specified by a comma — is only given to the leftmost piece of the command; the rest of the pieces have dot set for them by the execution of the pieces to their left.

As another simple example, consider a problem solved above: printing all lines in the file containing the word **Emacs**:

[†] The obvious analogy with shell pipelines is only partially valid, because the individual `sam` commands are all working on the same text; it is only how the text is sliced up that is changing.

```
,x/*\n/ g/Emacs/p
```

general introduction to the commands in Emacs and to try to show the method in the madness that is the Emacs command structure.

This command says to break the file into lines (`,x/*\n/`), and for each line that contains the string **Emacs** (`g/Emacs/`), run the command **p** with dot set to the line (not the match of **Emacs**), which prints the line. To save typing, because `.*\n` is a common pattern in x commands, if the x is followed immediately by a space, the pattern `.*\n` is assumed. Therefore, the above could be written more succinctly:

```
,x g/Emacs/p
```

The solution we used before was

```
,x/Emacs/+-p
```

which runs the command `+-p` with dot set to each match of **Emacs** in the file (recall that the idiom `+-p` prints the line containing the end of dot).

The two commands usually produce the same result (the `+-p` form will print a line twice if it contains **Emacs** twice). Which is better? `,x/Emacs/+-p` is easier to type and will be much faster if the file is large and there are few occurrences of the string, but it is really an odd special case. `,x/*\n/ g/Emacs/p` is slower — it breaks each line out separately, then examines it for a match — but is conceptually cleaner, and generalizes more easily. For example, consider the following piece of the Emacs manual:

```
command name="append-to-file", key="[unbound]"
Takes the contents of the current buffer and appends it to the
named file. If the file doesn't exist, it will be created.
```

```
command name="apropos", key="ESC-?"
Prompts for a keyword and then prints a list of those commands
whose short description contains that keyword. For example,
if you forget which commands deal with windows, just type
"@b[ESC-?}@t[window]@b[ESC-?]"
```

and so on

This text consists of groups of non-empty lines, with a simple format for the text within each group. Imagine that we wanted to find the description of the ‘apropos’ command. The problem is to break the file into individual descriptions, and then to find the description of ‘apropos’ and to print it. The solution is straightforward:

```
,x/(.+\\n)+/ g/command name="apropos"/p
command name="apropos", key="ESC-?"
Prompts for a keyword and then prints a list of those commands
whose short description contains that keyword. For example,
if you forget which commands deal with windows, just type
"@b[ESC-?}@t[window]@b[ESC-?]"
```

The regular expression `(.+\\n)+` matches one or more lines with one or more characters each, that is, the text between blank lines, so `,x/(.+\\n)+/` extracts each description; then `g/command name="apropos"/` selects the description for ‘apropos’ and **p** prints it.

Imagine that we had a C program containing the variable **n**, but we wanted to change it to **num**. This command is a first cut:

```
,x/n/ c/num/
```

but is obviously flawed: it will change all **n**’s in the file, not just the *identifier* **n**. A better solution is to use an x command to extract the identifiers, and then use **g** to find the **n**’s:


```
,x/[a-zA-Z][a-zA-Z_0-9]*/ g/n/ v!..! c/num/
```

It looks awful, but it's fairly easy to understand when read left to right. A C identifier is an alphabetic or underscore followed by zero or more alphanumerics or underscores, that is, matches of the regular expression `[a-zA-Z][a-zA-Z_0-9]*`. The `g` command selects those identifiers containing `n`, and the `v` is a trick: it rejects those identifiers containing more than one character. Hence the `c/num/` applies only to free-standing `n`'s.

There is still a problem here: we don't want to change `n`'s that are part of the character constant `\n`. There is a command `y`, complementary to `x`, that is just what we need: `y/pattern/command` runs the command on the pieces of text *between* matches of the pattern; if `x` selects, `y` rejects. Here is the final command:

```
,y/\n/ x/[a-zA-Z][a-zA-Z_0-9]*/ g/n/ v!..! c/num/
```

The `y/\n/` (with backslash doubled to make it a literal character) removes the two-character sequence `\n` from consideration, so the rest of the command will not touch it. There is more we could do here; for example, another `y` could be prefixed to protect comments in the code. I won't elaborate the example any further, but you should have an idea of the way in which the looping and conditional commands in **sam** may be composed to do interesting things.

Grouping

There is another way to arrange commands. By enclosing them in brace brackets `{}`, commands may be applied in parallel. This example uses the `=` command, which reports the line and character numbers of dot, together with `p`, to report on appearances of **Emacs** in our original file:

```
,p
This manual is organized in a rather haphazard manner. The first
several sections were written hastily in an attempt to provide a
general introduction to the commands in Emacs and to try to show
the method in the madness that is the Emacs command structure.
,x/Emacs/{
    =
    +-p
}
3; #171,#176
general introduction to the commands in Emacs and to try to show
4; #234,#239
the method in the madness that is the Emacs command structure.
```

(The number before the semicolon is the line number; the numbers beginning with `#` are character numbers.) As a more interesting example, consider changing all occurrences of **Emacs** to **vi** and vice versa. We can type

```
,x/Emacs\vi/{
    g/Emacs/ c/vi/
    g/vi/ c/Emacs/
}
```

or even

```
,x/[a-zA-Z]+/{
    g/Emacs/ v!.....! c/vi/
    g/vi/ v!...! c/Emacs/
}
```

to make sure we don't change strings embedded in words.

Multiple Changes

You might wonder why, once **Emacs** has been changed to **vi** in the above example, the second command in the braces doesn't put it back again. The reason is that the commands are run in parallel: within any top-level **sam** command, all changes to the file refer to the state of the file before any of the changes in that command are made. After all the changes have been determined, they are all applied simultaneously.

This means, as mentioned, that commands within a compound command see the state of the file before any of the changes apply. This method of evaluation makes some things easier (such as the exchange of **Emacs** and **vi**), and some things harder. For instance, it is impossible to use a **p** command to print the changes as they happen, because they haven't happened when the **p** is executed. An indirect ramification is that changes must occur in forward order through the file, and must not overlap.

Unix

sam has a few commands to connect to Unix processes. The simplest is **!**, which runs the command with input and output connected to the terminal.

```
!date
Wed May 28 23:25:21 EDT 1986
!
```

(When downloaded, the input is connected to **/dev/null** and only the first few lines of output are printed; any overflow is stored in **\$HOME/sam.err**.) The final **!** is a prompt to indicate when the command completes.

Slightly more interesting is **>**, which provides the current text as standard input to the Unix command:

```
1,2 >wc
2      22    131
!
```

The complement of **>** is, naturally, **<**: it replaces the current text with the standard output of the Unix command:

```
! <date
!
!p
Wed May 28 23:26:44 EDT 1986
```

The last command is **|**, which is a combination of **<** and **>**: the current text is provided as standard input to the Unix command, and the Unix command's standard output is collected and used to replace the original text. For example,

```
,| sort
```

runs **sort(1)** on the file, sorting the lines of the text lexicographically. Note that **<**, **>** and **|** are **sam** commands, not Unix shell operators.

The next example converts all appearances of **Emacs** to upper case using **tr(1)**:

```
,x/Emacs/ | tr a-z A-Z
```

tr is run once for each occurrence of **Emacs**. Of course, you could do this example more efficiently with a simple **c** command, but here's a trickier one: given a Unix mail box as input, convert all the **Subject** headers to distinct fortunes:

```
,x/^Subject:.*\n/ x/[^\n:]*\n/ < /usr/games/fortune
```

(The regular expression **[^\n:]** refers to any character *except* **:** and newline; the negation operator **^** excludes newline from the list of characters.) Again, **/usr/games/fortune** is run once for each **Subject** line, so each **Subject** line is changed to a different fortune.

A few other text commands

For completeness, I should mention three other commands that manipulate text. The **m** command moves the current text to after the text specified by the (obligatory) address after the command. Thus

```
/Emacs/+- m 0
```

moves the next line containing **Emacs** to the beginning of the file. Similarly, **t** (another historic character) copies the text:

```
/Emacs/+- t 0
```

would make, at the beginning of the file, a copy of the next line containing **Emacs**.

The third command is more interesting: it makes substitutions. Its syntax is *s/pattern/replacement/*. Within the current text, it finds the first occurrence of the pattern and replaces it by the replacement text, leaving dot set to the entire address of the substitution.

```
lp
```

This manual is organized in a rather haphazard manner. The first

```
s/haphazard/thoughtless/
```

```
p
```

This manual is organized in a rather thoughtless manner. The first

Occurrences of the character **&** in the replacement text stand for the text matching the pattern.

```
s/T/"&&&&"/
```

```
p
```

"TTTT"his manual is organized in a rather thoughtless manner. The first

There are two variants. The first is that a number may be specified after the **s**, to indicate which occurrence of the pattern to substitute; the default is the first.

```
s2/is/was/
```

```
p
```

"TTTT"his manual was organized in a rather thoughtless manner. The first

The second is that suffixing a **g** (global) causes replacement of all occurrences, not just the first.

```
s/[a-zA-Z]/x/g
```

```
p
```

"xxxx"xxx xxxxxx xxx xxxxxxxxxx xx x xxxxxx xxxxxxxxxxxx xxxxxxxx xxx xxxxx

Notice that in all these examples dot is left set to the entire line.

[The substitute command is vital to **ed**, because it is the only way to make changes within a line. It is less valuable in **sam**, in which the concept of a line is much less important. For example, many **ed** substitution idioms are handled well by **sam**'s basic commands. Consider the commands

```
s/good/bad/
```

```
s/good//
```

```
s/good/& bye/
```

which are equivalent in **sam** to

```
/good/c/bad/
```

```
/good/d
```

```
/good/a/ bye/
```

and for which the context search is likely unnecessary because the desired text is already dot. Also, beware this **ed** idiom:

```
1,$s/good/bad/
```

which changes the first **good** on each line; the same command in **sam** will only change the first one in the whole file. The correct **sam** version is

```
,x s/good/bad/
```

but what is more likely meant is

```
,x/good/ c/bad/
```

sam operates under different rules.]

Files

So far, we have only been working with a single file, but **sam** is a multi-file editor. Only one file may be edited at a time, but it is easy to change which file is the 'current' file for editing. To see how to do this, we need a **sam** with a few files; the easiest way to do this is to start it with a list of Unix file names to edit.

```
$ echo *.ms
conquest.ms death.ms emacs.ms famine.ms slaughter.ms
$ sam -d *.ms
-. conquest.ms
```

(I'm sorry the Horsemen don't appear in liturgical order.) The line printed by **sam** is an indication that the Unix file **conquest.ms** has been read, and is now the current file. **sam** does not read the Unix file until the associated **sam** file becomes current.

The **n** command prints the names of all the files:

```
n
-. conquest.ms
- death.ms
- emacs.ms
- famine.ms
- slaughter.ms
```

This list is also available in the menu on mouse button 3. The command **f** tells the name of just the current file:

```
f
-. conquest.ms
```

The characters to the left of the file name encode helpful information about the file. The minus sign becomes a plus sign if the file has a window open, and an asterisk if more than one is open. The period (another meaning of dot) identifies the current file. The leading blank changes to an apostrophe if the file is different from the contents of the associated Unix file, as far as **sam** knows. This becomes evident if we make a change.

```
ld
f
'-. conquest.ms
```

If the file is restored by an undo command, the apostrophe disappears.

```
u
f
-. conquest.ms
```

The file name may be changed by providing a new name with the **f** command:

```
f pestilence.ms
'-. pestilence.ms
```

f prints the new status of the file, that is, it changes the name if one is provided, and prints the name regardless. A file name change may also be undone.

```
u
f
-. conquest.ms
```

When **sam** is downloaded, the current file may be changed simply by selecting the desired file from the menu (selecting the same file subsequently cycles through the windows opened on the file). Otherwise, the **b** command can be used to choose the desired file:†

```
b emacs.ms
-. emacs.ms
```

Again, **sam** prints the name (actually, executes an implicit **f** command) because the Unix file **emacs.ms** is being read for the first time. It is an error to ask for a file **sam** doesn't know about, but the **B** command will prime **sam**'s menu with a new file, and make it current.

```
b flood.pic
?no such file 'flood.pic'
B flood.pic
-. flood.pic
n
- conquest.ms
- death.ms
- emacs.ms
- famine.ms
-. flood.pic
- slaughter.ms
```

Both **b** and **B** will accept a list of file names. **b** simply takes the first file in the list, but **B** loads them all. The list may be typed on one line —

```
B devil.tex satan.tex 666.tex emacs.tex
```

— or generated by a Unix command —

```
B <echo *.tex
```

The latter form requires a Unix command; **sam** does not understand the shell file name metacharacters, so **B *.tex** attempts to load a single file named ***.tex**. (The **<** form is of course derived from **sam**'s **<** command.) **echo** is not the only useful command to run subservient to **B**; for example,

```
B <grep -l Emacs *
```

will load only those files containing the string **Emacs**. Finally, a special case: a **B** with no arguments creates an empty, nameless file within **sam**.

The complement of **B** is **D**:

```
D devil.tex satan.tex 666.tex emacs.tex
```

eradicates the files from **sam**'s memory (not from the Unix machine's disc). **D** without any file names removes the current file from **sam**.

There are three other commands that relate the current file to Unix files. The **w** command writes the file to disc; without arguments, it writes the entire file to the Unix file associated with the current file in **sam** (it is the only command whose default address is not dot). Of course, you can specify an address to be written, and a different file name, with the obvious syntax:

† A bug prevents the **b** command from working when downloaded. Because the menu is more convenient anyway, and because the method of choosing files from the command language is slated to change, the bug hasn't been fixed.

```
1,2w /tmp/revelations
/tmp/revelations: #44
```

sam responds with the file name and the number of characters written to the file. The **write** command on the button 3 menu is identical in function to an unadorned **w** command.

The other two commands, **e** and **r**, read data from Unix files. The **e** command clears out the current file, reads the data from the named file (or uses the current file's old name if none is explicitly provided), and sets the file name. It's much like a **B** command, but puts the information in the current file instead of a new one. **e** without any file name is therefore an easy way to refresh **sam**'s copy of a Unix file. [Unlike in **ed**, **e** doesn't complain if the file is modified. The principle is not to protect against things that can be undone if wrong.] Since its job is to replace the whole text, **e** never takes an address.

The **r** command is like **e**, but it doesn't clear the file: the text in the Unix file replaces dot, or the specified text if an address is given.

```
r emacs.ms
```

has essentially the effect of

```
<cat emacs.ms
```

The commands **r** and **w** will set the name of the file if the current file has no name already defined; **e** sets the name even if the file already has one.

There is a command, analogous to **x**, that iterates over files instead of pieces of text: **X** (capital **x**). The syntax is easy; it's just like that of **x** — *X/pattern/command*. (The complementary command is **Y**, analogous to **y**.) The effect is to run the command in each file whose menu entry (that is, whose line printed by an **f** command) matches the pattern. For example, since an apostrophe identifies modified files,

```
X/' w
```

writes the changed files out to disc. Here is a longer example: find all uses of a particular variable in the C source files:

```
X\c$/ ,x/variable/+-p
```

We can use an **f** command to identify which file the variable appears in:

```
X\c$/ ,g/variable/ {
  f
  ,x/variable/+-{
    =
    p
  }
}
```

Here, the **g** command guarantees that only the names of files containing the variable will be printed (but beware that **sam** may confuse matters by printing the names of files it reads in during the command). The **=** command shows where in the file the variable appears, and the **p** command prints the line.

The **D** command is handy as the target of an **X**. This example deletes from the menu all C files that do not contain a particular variable:

```
X\c$/ ,v/variable/ D
```

If no pattern is provided for the **X**, the command (which defaults to **f**) is run in all files, so

```
X D
```

cleans **sam** up for a fresh start.

But rather than working any further, let's stop now:

q
\$

Some of the file manipulating commands can be undone: undoing a **f**, **e**, or **r** restores the previous state of the file, but **w**, **B** and **D** are irrevocable. And, of course, so is **q**.

NAME

sam – screen editor with structural regular expressions

SYNOPSIS

sam [*option* ...] [*files*]

sam -r *machine*

sam.save

B [**-nnnn**] *file* ...

DESCRIPTION

Sam is a multi-file editor. It modifies a local copy of an external file. The copy is here called a *file*. The files are listed in a menu available through mouse button 3 or the **n** command. Each file has an associated name, usually the name of the external file from which it was read, and a ‘modified’ bit that indicates whether the editor’s file agrees with the external file. The external file is not read into the editor’s file until it first becomes the current file—that to which editing commands apply—whereupon its menu entry is printed. The options are 289.sp 28u

-d Do not download the terminal part of *sam*. Editing will be done with the command language only, as in *ed*(1).

-r *machine*

Run the host part remotely on the specified machine, the terminal part locally.

-s *file* Start the host part from the indicated file on the remote host. Only meaningful with the **-r** option.

-t *path* Start the terminal part from the indicated file. Useful for debugging.

The standard X11 **-geom** toolkit option can be used to select the desired window geometry.

Regular expressions

Regular expressions are as in *regex*(6) with the addition of **\n** to represent newlines. A regular expression may never contain a literal newline character. The elements of regular expressions are: 384.sp 0u

. Match any character except newline.

\n Match newline.

\x For any character except **n** match the character (here *x*).

[abc] Match any character in the square brackets. **\n** may be mentioned.

[^abc] Match any character not in the square brackets, but never a newline. Both these forms accept a range of ASCII characters indicated by a dash, as in **a-z**.

^ Match the null string immediately after a newline.

\$ Match the null string immediately before a newline.

Any other character except newline matches itself.

In the following, *r1* and *r2* are regular expressions. 359.sp 0u

(r1) Match what *r1* matches.

r1/r2 Match what *r1* or what *r2* matches.

r1* Match zero or more adjacent matches of *r1*.

r1+ Match one or more adjacent matches of *r1*.

r1? Match zero or one matches of *r1*.

The operators *****, **+** and **?** are highest precedence, then catenation, then **|** is lowest. The empty regular expression stands for the last complete expression encountered. A regular expression in *sam* matches the longest leftmost substring formally matched by the expression. Searching in the reverse direction is equivalent to searching backwards with the catenation operations reversed in the expression.

Addresses

An address identifies a substring in a file. In the following, ‘character *n*’ means the null string after the *n*-th character in the file, with 1 the first character in the file. ‘Line *n*’ means the *n*-th match, starting at the beginning of the file, of the regular expression **.*\n?**. (The peculiar properties of a last line without a newline are temporarily undefined.) All files always have a current substring, called dot, that

is the default address.

Simple Addresses

- #*n*** The empty string after character *n*; **#0** is the beginning of the file.
- n*** Line *n*.
- /regex/***
- ?*regex*?** The substring that matches the regular expression, found by looking toward the end (*/*) or beginning (*?*) of the file, and if necessary continuing the search from the other end to the starting point of the search. The matched substring may straddle the starting point.
- 0** The string before the first full line. This is not necessarily the null string; see **+** and **-** below.
- \$** The null string at the end of the file.
- .** Dot.
- '** The mark in the file (see the **k** command below).
- "*regex*"** Preceding a simple address (default **.**), refers to the address evaluated in the unique file whose menu line matches the regular expression.

Compound Addresses

In the following, *a1* and *a2* are addresses. 387.sp 28u

- a1+a2*** The address *a2* evaluated starting at the end of *a1*.
 - a1-a2*** The address *a2* evaluated looking in the reverse direction starting at the beginning of *a1*.
 - a1,a2*** The substring from the beginning of *a1* to the end of *a2*. If *a1* is missing, **0** is substituted. If *a2* is missing, **\$** is substituted.
 - a1;a2*** Like *a1,a2*, but with *a2* evaluated at the end of, and dot set to, *a1*.
- The operators **+** and **-** are high precedence, while **,** and **;** are low precedence. In both **+** and **-** forms, if *a2* is a line or character address with a missing number, the number defaults to 1. If *a1* is missing, **.** is substituted. If both *a1* and *a2* are present and distinguishable, **+** may be elided. *a2* may be a regular expression; if it is delimited by **?'s**, the effect of the **+** or **-** is reversed. It is an error for a compound address to represent a malformed substring. Some useful idioms: *a1+-* (*a1-+*) selects the line containing the end (beginning) of *a1*. **0/*regex*/** locates the first match of the expression in the file. (The form **0//** sets dot unnecessarily.) ***/regex*///** finds the second following occurrence of the expression, and ***./regex/*** extends dot.

Commands

In the following, text demarcated by slashes represents text delimited by any printable ASCII character except alphanumerics. Any number of trailing delimiters may be elided, with multiple elisions then representing null strings, but the first delimiter must always be present. In any delimited text, newline may not appear literally; **\n** may be typed for newline; and **\V** quotes the delimiter, here **/**. Backslash is otherwise interpreted literally, except in **s** commands.

Most commands may be prefixed by an address to indicate their range of operation. Those that may not are marked with a ***** below. If a command takes an address and none is supplied, dot is used. The sole exception is the **w** command, which defaults to **0,\$**. In the description, 'range' is used to represent whatever address is supplied. Many commands set the value of dot as a side effect. If so, it is always set to the 'result' of the change: the empty string for a deletion, the new text for an insertion, etc. (but see the **s** and **e** commands).

Text commands**a***text/*

or

a*lines of text***.** Insert the text into the file after the range. Set dot.**c****i** Same as **a**, but **c** replaces the text, while **i** inserts *before* the range.**d** Delete the text in the range. Set dot.**s***/regexp/text/*

Substitute *text* for the first match to the regular expression in the range. Set dot to the modified range. In *text* the character **&** stands for the string that matched the expression. Backslash behaves as usual unless followed by a digit: **\d** stands for the string that matched the subexpression begun by the *d*-th left parenthesis. If *s* is followed immediately by a number *n*, as in **s2/x/y/**, the *n*-th match in the range is substituted. If the command is followed by a **g**, as in **s/x/y/g**, all matches in the range are substituted.

m *al***t** *al* Move the range to after *al* (**m**), or copy it (**t**). Set dot.**Display commands****p** Print the text in the range. Set dot.**=** Print the line address and character address of the range.**=#** Print just the character address of the range.**File commands***** b** *file-list*

Set the current file to the first file named in the list that *sam* also has in its menu. The list may be expressed *<shell command* in which case the file names are taken as words (in the shell sense) generated by the shell command.

*** B** *file-list*

Same as **b**, except that file names not in the menu are entered there, and all file names in the list are examined.

*** n** Print a menu of files. The format is:

' or blank	indicating the file is modified or clean,
- or +	indicating the file is unread or has been read (in the terminal, * means more than one window is open),
. or blank	indicating the current file,
a blank,	
	and the file name.

*** D** *file-list*

Delete the named files from the menu. If no files are named, the current file is deleted. It is an error to **D** a modified file, but a subsequent **D** will delete such a file.

I/O Commands*** e** *filename*

Replace the file by the contents of the named external file. Set dot to the beginning of the file.

r *filename*

Replace the text in the range by the contents of the named external file. Set dot.

w *filename*

Write the range (default 0,\$) to the named external file.

*** f** *filename*

Set the file name and print the resulting menu entry.

If the file name is absent from any of these, the current file name is used. **e** always sets the file name, **r** and **w** do so if the file has no name.

< *shell-command*

Replace the range by the standard output of the shell command.

> *shell-command*

Sends the range to the standard input of the shell command.

| *shell-command*

Send the range to the standard input, and replace it by the standard output, of the shell command.

* ! *shell-command*

Run the shell command.

* **cd** *directory*

Change working directory. If no directory is specified, **\$HOME** is used.

In any of <, >, | or !, if the *shell command* is omitted the last *shell command* (of any type) is substituted. If *sam* is downloaded, ! sets standard input to **/dev/null**, and otherwise unassigned output (**stdout** for ! and >, **stderr** for all) is placed in **\$HOME/sam.err** and the first few lines are printed.

Loops and Conditionals

x/regex/ command

For each match of the regular expression in the range, run the command with dot set to the match. Set dot to the last match. If the regular expression and its slashes are omitted, **/.*\n/** is assumed. Null string matches potentially occur before every character of the range and at the end of the range.

y/regex/ command

Like **x**, but run the command for each substring that lies before, between, or after the matches that would be generated by **x**. There is no default behavior. Null substrings potentially occur before every character in the range.

* *X/regex/ command*

For each file whose menu entry matches the regular expression, make that the current file and run the command. If the expression is omitted, the command is run in every file.

* *Y/regex/ command*

Same as **X**, but for files that do not match the regular expression, and the expression is required.

g/regex/ command

v/regex/ command

If the range contains (**g**) or does not contain (**v**) a match for the expression, set dot to the range and run the command.

These may be nested arbitrarily deeply, but only one instance of either **X** or **Y** may appear in a single command. An empty command in an **x** or **y** defaults to **p**; an empty command in **X** or **Y** defaults to **f**. **g** and **v** do not have defaults.

Miscellany

k Set the current file's mark to the range. Does not set dot.

* **q** Quit. It is an error to quit with modified files, but a second **q** will succeed.

* **u** *n* Undo the last *n* (default 1) top-level commands that changed the contents or name of the current file, and any other file whose most recent change was simultaneous with the current file's change. Successive **u**'s move further back in time. The only commands for which **u** is ineffective are **cd**, **u**, **q**, **w** and **D**.

(empty) If the range is explicit, set dot to the range. If *sam* is downloaded, the resulting dot is selected on the screen; otherwise it is printed. If no address is specified (the command is a newline) dot is extended in either direction to line boundaries and printed. If dot is thereby unchanged, it is set to **.+1** and printed.

Grouping and multiple changes

Commands may be grouped by enclosing them in braces **{}**. Commands within the braces must appear on separate lines (no backslashes are required between commands). Semantically, an opening brace is like a command: it takes an (optional) address and sets dot for each sub-command. Commands within the braces are executed sequentially, but changes made by one command are not visible to other commands (see the next paragraph). Braces may be nested arbitrarily.

When a command makes a number of changes to a file, as in `x/re/c/text/`, the addresses of all changes to the file are computed in the original file. If the changes are in sequence, they are applied to the file. Successive insertions at the same address are catenated into a single insertion composed of the several insertions in the order applied.

The terminal

What follows refers to behavior of *sam* when downloaded, that is, when operating as a display editor on a bitmap display. This is the default behavior; invoking *sam* with the `-d` (no download) option provides access to the command language only.

Each file may have zero or more windows open. Each window is equivalent and is updated simultaneously with changes in other windows on the same file. Each window has an independent value of dot, indicated by a highlighted substring on the display. Dot may be in a region not within the window. There is usually a 'current window', marked with a dark border, to which typed text and editing commands apply. The escape key (ESC) selects (sets dot to) text typed since the last mouse button hit.

The button 3 menu controls window operations. The top of the menu provides the following operators, each of which uses one or more cursors to prompt for selection of a window or sweeping of a rectangle.

- new** Create a new, empty file: Depress button 3 where one corner of the new rectangle should appear (box cursor), and move the mouse while holding down button 3 to the diagonally opposite corner. 'Sweeping' a null rectangle gets a large window, disjoint from the command window or the whole sam window, depending on where the null rectangle is.
- xerox** Create a copy of an existing window. After selecting the window to copy with button 1, sweep out the window for the copy.
- reshape** Change the size and location of a window. First click button 3 in the window to be changed (gunsight cursor). Then sweep out a window as for the **new** menu selection.
- close** Delete the window. In the last window of a file, **close** is equivalent to a **D** for the file.
- write** Equivalent to a **w** for the file.

Below these operators is a list of available files, starting with `~~sam~~`, the command window. Selecting a file from the list makes the most recently used window on that file current, unless it is already current, in which case selections cycle through the open windows. If no windows are open on the file, the user is prompted to open one. Files other than `~~sam~~` are marked with one of the characters `-+*` according as zero, one, or more windows are open on the file. A further mark `.` appears on the file in the current window and a single quote, `'`, on a file modified since last write.

The command window, created automatically when *sam* starts, is an ordinary window except that text typed to it is interpreted as commands for the editor rather than passive text, and text printed by editor commands appears in it. There is an 'output point' that separates commands being typed from previous output. Commands typed in the command window apply to the current open file—the file in the most recently current window.

Manipulating text

Typed characters replace the current selection (dot) in the current window. Backspace deletes the previous character. Escape selects (sets dot to) everything typed since the last mouse hit.

Button 1 changes selection. Pointing to a non-current window with button 1 makes it current; within the current window, button 1 selects text, thus setting dot. Double-clicking selects text to the boundaries of words, lines, quoted strings or bracketed strings, depending on the text at the click.

Button 2 provides a menu of editing commands:

- cut** Delete dot and save the deleted text in the snarf buffer.
- paste** Replace the text in dot by the contents of the snarf buffer.
- snarf** Save the text in dot in the snarf buffer.

- look** Search forward for the next occurrence of the literal text in dot. If dot is the null string, the text in the snarf buffer is used. The snarf buffer is unaffected.
- <exch>** Exchange the snarf buffer with the current selection in another X11 window. The exchange of a large amount of selected text is truncated to the size of Sam's internal snarf buffer (currently 4K) without warning.
- /regexp** Search forward for the next match of the last regular expression typed in a command. (Not in command window.)
- send** Send the text in dot, or the snarf buffer if dot is the null string, as if it were typed to the command window. Saves the sent text in the snarf buffer. (Command window only.)

X11 resources

Various attributes of *sam* can be set by giving values to X11 resources for the class **Sam**. There are various ways to do this; one is to have a file called **Sam** in your home directory, with entries in it like:

```
Sam*width: 500
Sam*height: 600
Sam*font: fixed
Sam*scrollForwardR: true
```

In addition to the usual X11 toolkit resources, the **scrollForwardR** resource says where the right button (button 3) scrolls forward or backward in the file.

Abnormal termination

If *sam* terminates other than by a **q** command (by hangup, deleting its window, etc.), modified files are saved in an executable file, **\$HOME/sam.save**. This program, when executed, asks whether to write each file back to an external file. The answer **y** causes writing; anything else skips the file. If a machine crash prevents the creation of a **sam.save** file, all changes are lost. If an editing session is difficult to replicate, you should write your changed files often.

B is a shell-level command that causes an instance of *sam* running on the same terminal to load the named *files*. The option allows a line number to be specified for the initial position to display in the last named file.

FILES

```
$HOME/sam.save
$HOME/sam.err
<system-dependent directory>/samsave the program called to unpack $home/sam.save.
```

SEE ALSO

ed(1), *regexp*(6)

BUGS

When a *sam* window is resized, the command window may have the wrong size.